

WHAT IS CLAIMED IS:

1. A method of sending a message in an interconnection fabric, wherein the interconnection fabric couples together a plurality of nodes, wherein each node of the plurality of nodes comprises a plurality of input ports and a plurality of output ports,
5 comprising:

identifying a route in the interconnection fabric for sending the message from a sending node to a destination node;

10

encoding a routing directive in the message, wherein the routing directive describes the route and comprises at least one segment, wherein each segment comprises a direction component and a distance component;

15

sending the message on one of the output ports of the sending node;

receiving the message on one of the input ports of a first node connected to the output port of the sending node;

20

decrementing the distance component for a current segment of the routing directive;

selecting one of the output ports of the first node according to the current segment of the routing directive in the message; and

25

sending the message on the selected one of the output ports of the first node.

30

2. The method as recited in claim 1, wherein said selecting comprises:

if, after said decrementing, the distance component for the current segment is greater than zero, selecting the output port corresponding to a same routing direction in which the message was traveling when received; and

if, after said decrementing, the distance component for the current segment is zero, selecting the output port corresponding to the direction component of the current segment.

3. The method as recited in claim 2, wherein if, after said decrementing, the distance component for the current segment is zero, and the output port is selected according to the direction component of the current segment, the method further comprises removing the current segment from the routing directive so that a next segment becomes the current segment when the message is sent on the selected output port.

4. The method as recited in claim 3, wherein the routing directive further comprises a pointer to the current segment, and wherein said removing the current segment comprises moving the pointer to the next segment.

5. The method as recited in claim 1, further comprising:

a subsequent node receiving the message;

the subsequent node decrementing the distance component for the current segment of the routing directive;

wherein after said decrementing:

if the distance component for the current segment is greater than zero, the subsequent node selecting the output port corresponding to a same routing direction in which the message was traveling when received; and

5

if the distance component for the current segment is zero, the subsequent node selecting a port corresponding to the direction component of the current segment.

10 6 The method as recited in claim 5, wherein the subsequent node selecting a port corresponding to the direction component comprises:

selecting the corresponding output port if the direction component for the current segment specifies a routing direction; and

15

selecting a device port if the direction component for the current segment specifies that the subsequent node is the destination for the message.

7. The method as recited in claim 1 wherein said identifying the route comprises
20 selecting the routing directive from a routing table.

8. The method as recited in claim 7 wherein the routing table comprises a plurality of independent routes from the sending node to the destination node.

25 9. The method as recited in claim 1 wherein the interconnection fabric is a torus interconnection fabric.

10. The method as recited in claim 1, further comprising:

30 identifying a return route from the destination node to the sending node; and

5 encoding a return routing directive in the message, wherein the return routing directive describes the return route and comprises at least one segment, wherein each segment comprises a direction component and a distance component.

11. The method as recited in claim 10, further comprising calculating the return routing directive.

10 12. The method as recited in claim 11, wherein the interconnection fabric is bi-directional, and wherein calculating the return routing directive comprises reversing the routing directive.

15 13. The method as recited in claim 1, further comprising incrementally encoding a return routing directive in the message, wherein the return routing directive describes a return route from the destination node to the sending node and comprises at least one segment, and wherein each segment comprises a direction component and a distance component.

20 14. The method as recited in claim 13, wherein incrementally encoding comprises:

incrementing the distance component for a current segment of the return routing directive;

25 wherein if, after said decrementing, the distance component for the current segment of the routing directive is zero, the method further comprises modifying the direction component of a current segment of the return routing directive and adding a new segment to the return routing directive so that the new segment becomes the current segment of the return routing
30 directive when the message is sent on the selected output port.

15. The method as recited in claim 14, wherein the return routing directive further comprises a pointer to the current segment, wherein adding a new segment to the return routing directive further comprises moving the pointer to the new segment.

5

16. The method as recited in claim 1 wherein a first number of segments of a first routing directive differs from a second number of segments of a second routing directive.

17. The method as recited in claim 3 further comprising a subsequent node receiving the message and, if all of the segments of the routing directive have been removed, the subsequent node identifying itself as the destination node and selecting a device port.

10

18. The method as recited in claim 1, wherein each direction component comprises a direction relative to a routing direction the message was traveling in when received.

15

19. The method as recited in claim 1, wherein each direction component comprises an identifier of one of the output ports of one of the nodes.

20

20. The method as recited in claim 1, wherein the destination node is configured to communicate with a storage device.

20

21. The method as recited in claim 20, wherein the storage device comprises a disk drive.

25

22. A node, comprising:

a routing unit;

a plurality of input ports; and

30

a plurality of output ports;

wherein the node is configured to be connected to an interconnection fabric,
wherein the interconnection fabric is configured to connect the node to a
plurality of nodes;

wherein the routing unit is configured to receive a message being sent along a
route from a sending node to a destination node in the interconnection
fabric;

wherein the routing unit is further configured to receive a routing directive
encoded in the message, wherein the routing directive describes the route
and comprises at least one segment, and wherein a segment comprises a
direction component and a distance component; and

wherein the node is configured to receive the message on one of the input ports
when the node is not the sending node, wherein the node is further
configured to decrement the distance component of a current segment of
the routing directive and to select one of the output ports according to the
current segment.

23. The node as recited in claim 22, wherein the routing unit is configured to identify
the route and to encode the routing directive in the message when the node is the sending
node, and wherein the node is configured to send the message on one of the output ports.

24. The node as recited in claim 22, wherein the node is configured to communicate
with a device on a device port, wherein the device is configured to select a route, encode
a routing directive in the message and communicate the message to the node on the
device port when the node is the sending node.

25. The node as recited in claim 24, wherein the node is further configured to select one of the output ports according to the current segment.

26. The node as recited in claim 22, wherein the node is configured to select:

5

one of the output ports corresponding to a same routing direction in which the message was traveling when received if, after said decrementing, the distance component for the current segment is greater than zero; and

10 one of the output ports corresponding to the direction component of the current segment if, after said decrementing, the distance component for the current segment is zero.

15 27. The node as recited in claim 26, wherein the node is further configured to remove the current segment from the routing directive if, after said decrementing, the distance component for the current segment is zero, and the wherein the node is configured to select the output port according to the direction component of the current segment, so that a next segment becomes the current segment when the message is sent on the selected output port.

20

28. The node as recited in claim 27, wherein the routing directive further comprises a pointer to the current segment, and wherein said being configured to remove the current segment comprises being configured to move the pointer to the next segment.

25 29. The node as recited in claim 22, wherein the node is configure to select:

one of the output ports corresponding to a same routing direction in which the message was traveling when received if, after said decrementing, the distance component for the current segment is greater than zero;

30

one of the output ports corresponding to the direction component of the current segment if, after said decrementing, the distance component for the current segment is zero, and if the direction component for the current segment does not identify that the node is the destination node; and

5

a device port if, after said decrementing, the distance component for the current segment is zero and if the direction component for the current segment identifies that the node is the destination node.

10 30. The node as recited in claim 22, wherein the routing unit is configured to identify the route by selecting the routing directive from a routing table.

31. The node as recited in claim 30, wherein the routing table comprises a plurality of independent routes from the sending node to the destination node.

15

32. The node as recited in claim 22, wherein the interconnection fabric comprises a torus interconnection fabric.

20

33. The node as recited in claim 22, wherein, if the node is the sending node, the routing unit is further configured to identify a return route from the destination node to the sending node and to encode a return routing directive in the message, wherein the return routing directive describes the return route and comprises at least one segment, wherein each segment comprises a direction component and a distance component.

25

34. The node as recited in claim 33, wherein, if the node is the sending node, the routing unit is further configured to calculate the return routing directive.

30

35. The node as recited in claim 34, wherein the interconnection fabric is bi-directional, and wherein calculating the return routing directive comprises reversing the routing directive.

36. The node as recited in claim 22, wherein the node is configured to communicate with a RAID controller.

5 37. The node as recited in claim 22, wherein the node is configured to communicate with a mass storage device.

38. The node as recited in claim 37, wherein the mass storage device is a disk drive.

10 39. A device, comprising:

an interface configured to communicate with a source node in an interconnection fabric, wherein the interconnection fabric comprises a plurality of routes between the source node and a destination node; and

15

a controller configured to provide a first routing directive describing a first route from the source node to the destination node, wherein the routing directive comprises at least one segment, wherein each segment comprises a distance component and a direction component, wherein the distance component is configured to be decremented by a receiving node;

20

and wherein the controller is further configured to encode the first routing directive in a message, and to communicate the message to the source node to be sent on the interconnection fabric to the destination node.

25

40. The device of claim 39, wherein said controller comprises a RAID controller.

41. The device of claim 39, wherein the controller comprises a host interface configured to communicate with a host computer.

30

42. The device of claim 39, wherein the controller comprises a disk storage device controller.

43. The device of claim 39, wherein the controller is further configured to maintain a routing table comprising the first routing directive, and wherein the controller is further configured to select the first routing directive from the routing table.

44. The device of claim 43, wherein the routing table further comprises a second routing directive describing a second route from the source node to the destination node.

10

45. The device of claim 44, wherein the second routing directive comprises a different number of segments than the first routing directive.

46. The device of claim 39, wherein the controller is further configured to calculate the first routing directive.

47. The device of claim 39, wherein the controller is further configured to provide a return routing directive describing a return route from the destination node to the source node, and wherein the controller is further configured to encode the return routing directive in the message.

20

48. The device of claim 47, wherein the controller is further configured to select the return routing directive from a routing table.

49. The device of claim 47, wherein the controller is further configured to calculate the return routing directive from the first routing directive.

25

50. The device of claim 39, wherein the controller is further configured to encode a return routing directive describing a return route from the destination node to the source

node in the message, and wherein the return routing directive is configured to be incrementally added to as the message is routed to the destination node.

51. The device of claim 50, wherein the return routing directive is further configured
5 to be used to return an error message to the source node if a routing error is encountered.

52. The device of claim 51, wherein the controller is further configured to use the
incrementally created return routing directive to locate the routing error if an error
message is returned, wherein the incrementally created return routing directive indicates a
10 last node that successfully received the message.

CONFIDENTIAL